

Statistical Process Control Using Dynamic Sampling Scheme

Zhonghua Li¹ and Peihua Qiu²

¹LPMC and School of Mathematical Sciences, Nankai University, Tianjin 300071, China

²Department of Biostatistics, University of Florida, Gainesville, FL 32610, USA

Abstract

This paper considers statistical process control (SPC) of univariate processes, and tries to make two contributions to the univariate SPC problem. First, we propose a continuously variable sampling scheme, based on a quantitative measure of the likelihood of a process distributional shift at each observation time point, provided by the p -value of the conventional cumulative sum (CUSUM) charting statistic. For convenience of the design and implementation, the variable sampling scheme is described by a parametric function in the flexible Box-Cox transformation family. Second, the resulting CUSUM chart using the variable sampling scheme is combined with an adaptive estimation procedure for determining its reference value, to effectively protect against a range of unknown shifts. Numerical studies show that it performs well in various cases. A real data example from a chemical process illustrates the application and implementation of our proposed method. This paper has supplementary materials online.

Key words: Adaptive Estimation; Bootstrap; Dynamic Sampling; Monte Carlo Simulation; Variable Sampling.

1 Introduction

Statistical process control (SPC) charts provide us an effective tool for monitoring the performance of a sequential process over time. They have broad applications in manufacturing industries and in biology, genetics, medicine, finance and many other areas as well (Montgomery 2009). Three widely used control charts include the Shewhart chart (Shewhart 1931), cumulative sum (CUSUM) chart (Page 1954), and exponentially weighted moving average (EWMA) chart (Roberts 1959). See, e.g., Chen et al. (2001), Chen and Zhang (2005), Hawkins and

Olwell (1998), Montgomery (2007, 2009), Qiu (2013), Woodall (2000), and Yeh et al. (2004) for related discussions.

The three conventional control charts mentioned above have their own strengths. For instance, the Shewhart chart is more effective in detecting large and isolated shifts while the CUSUM and EWMA control charts are more efficient in detecting small and persistent shifts (Lucas and Saccucci 1990). These charts give a signal whenever their charting statistics fall outside their control limits. By using these charts, users would know whether the process is in-control (IC) or not at each time point. But, the charts cannot provide us a quantitative measure of the likelihood of a potential process distributional shift, even after a signal of shift has been given. In practice, such a quantitative measure is useful for taking appropriate subsequent actions. For instance, in cases when a variable sampling scheme is possible (Reynolds et al. 1990), even if a signal of shift is not delivered at a given time point, the quantitative measure of the shift likelihood can be used for adjusting the next sampling time properly. We can wait longer for the next sampling time if the shift likelihood is smaller, and shorter otherwise. In the context of hypothesis testing, the p -value approach is often preferred nowadays, compared to the conventional “rejection region” approach, because the former provides us not only a decision about whether the null hypothesis should be rejected at a given significance level but also a quantitative measure of the evidence in the observed data to support such a decision.

In recent years, control charts with variable sampling rate (VSR) schemes have attracted considerable attention of statisticians. As pointed out by Montgomery (2007), one important area of SPC research continues to be the use of adaptive control charts with variable sample sizes and/or sampling intervals. By a VSR scheme, the sampling rate changes over time, depending on the current and prior sampling results. One major advantage of a VSR chart, compared to a chart using a fixed sampling rate (FSR), is that it can detect small to moderate shifts effectively, given the IC average run length (denoted as ARL_0), and the IC average sampling rate. There are several different ways to change the sampling rate, including the variable sampling intervals (VSI), variable sample sizes (VSS), and variable sample sizes and sampling intervals (VSSI). See Montgomery (2009) for a more detailed description. In the literature (Costa 1998, Wu et al. 2007, Reynolds and Arnold 2001), the sampling interval function $d(\cdot)$ usually takes only two possible values. Recently, Stoumbos et al. (2011) derived the optimal sampling interval function in terms of the adjusted average time to signal (AATS) using a quadratic programming algorithm. They showed that this optimal function could be approximated by a function with

only two values d_1 and d_2 . Thus, they recommended using this simplified sampling interval function in practice. However, they pointed out that a control chart using such a simplified sampling interval function would be “strongly tuned towards the size of shift for which the charts have been optimized, at the significant expense of some of the other shift sizes.” Moreover, they found that “with a dual-sampling-interval policy, large values of d_2 make the chart more efficient for small shifts, and small values of d_2 make it more efficient for large shifts.” Therefore, although the simplified sampling interval function is recommended because of its simplicity, it is not optimal, and has much room for improvement.

In this paper, we suggest choosing $d(\cdot)$ to be a continuous function of the p -value of the charting statistic of our proposed CUSUM chart. The sampling scheme with such a sampling interval function is called a dynamic sampling scheme in this paper. For convenience of the design and implementation of our proposed control chart, we suggest approximating the sampling interval function by a parametric function in the flexible Box-Cox transformation family. It will be shown that such a dynamic sampling scheme has good performance in various cases.

It is well known that the CUSUM chart has an attractive theoretical property that it is the optimal shift detection procedure if its reference value, denoted by k , is chosen properly for a particular shift size (Lorden 1971, Pollack 1985, Moustakides 1986). Here, the optimality means that, among all control charts with their ARL_0 values equal or smaller than a given value, the CUSUM chart has the smallest value of $\max_{\tau} ARL_1$, where τ is the true position of a shift with a give size and ARL_1 is the out-of-control (OC) average run length of the CUSUM chart for detecting that shift. See Qiu (2013, Chapter 4) for a detailed discussion about the CUSUM chart. In practice, however, the size of a potential shift is unknown. To overcome this difficulty, for the VSR chart with a simplified sampling interval function, Stoumbos et al. (2011) recommended that d_1 be chosen as small as possible, d_2 chosen between 1.20 and 3.00 time units, and k chosen between 0.10 and 0.60. As mentioned above, such a recommended design is tuned towards certain shift sizes at the significant expense of some other shift sizes. To overcome this limitation, in this paper, we suggest using the scheme by Sparks (2000) to adaptively estimate the size of a potential shift at each time point and then choose the reference value of our proposed CUSUM chart accordingly.

As far as we know, this paper is the first to define the sampling interval function $d(\cdot)$ as a function of the p -value of the charting statistic of a control chart. Unlike most existing control charts with VSR schemes, in which $d(\cdot)$ usually takes only two values, we suggest adjusting $d(\cdot)$

by a continuous parametric function of the p -value. This paper is also the first to combine the idea to use a continuously-changing sampling interval function $d(\cdot)$ and the idea by Sparks (2000) to adaptively estimate the shift size. The remainder of the article is organized as follows. In the next section, our proposed CUSUM chart is described in detail. Its numerical performance is investigated in Section 3. Then, our method is demonstrated using a real data example in Section 4. Some remarks conclude the article in Section 5. Some numerical results, computer codes, and the real-data discussed in Section 4 are provided in online supplementary files.

2 Proposed CUSUM Chart

We describe our proposed CUSUM chart in three parts. In Section 2.1, we give a brief introduction to CUSUM charts using p -values. In Section 2.2, the dynamic sampling scheme is described. Then, in Section 2.3, adaptive selection of the reference value of the proposed CUSUM chart is discussed.

2.1 CUSUM charts using p -values

In the literature, there have been some discussions about designing a control chart using the p -value of its charting statistic. For instance, in their technical reports, Benjamini and Kling (1999, 2002, 2007) proposed the idea to present the p -values of the charting statistic of a traditional control chart (e.g., the \bar{X} -chart) for monitoring processes in cases when the IC process distribution was completely known. In their proposal, the p -value was computed using the Markov chain representation of the charting statistic (e.g., Brook and Evans 1972). In cases when the IC process distribution is assumed normal with a known variance, Grigg and Spiegelhalter (2008) provided an approximation formula for the IC distribution of the charting statistic of the conventional CUSUM chart. Li and Tsung (2009) studied the false discovery rate in multistage process monitoring, in which p -value calculation of the charting statistics used in different stages of process monitoring is discussed. Li et al. (2013) systematically discussed how to design control charts using p -values in cases when the IC process distribution was completely known, in cases when the IC process distribution had a parametric form with unknown parameters, and in cases when the IC process distribution was completely unknown. The last scenario (unknown IC process distribution) is briefly introduced below.

Let X_1, X_2, \dots, X_τ be a sequence of independent and identically distributed (i.i.d.) random variables with mean μ_0 , variance σ^2 , and a cumulative distribution function (cdf) $F_0(\cdot)$, and $X_{\tau+1}, X_{\tau+2}, \dots$ be a sequence of i.i.d. random variables with the same distribution except that their common mean is $\mu_1 \neq \mu_0$, where τ is an unknown change point. Then, the charting statistic of the conventional CUSUM for detecting an upward mean shift is defined by

$$\begin{cases} C_0^+ = 0, \\ C_n^+ = \max(0, C_{n-1}^+ + X_n - \mu_0 - k), \end{cases}$$

where $k > 0$ is its reference value. If k is chosen as $(\mu_1 - \mu_0)/2$, then the chart is optimal for detecting the particular shift μ_1 . The chart gives a signal of an upward mean shift when

$$C_n^+ > h,$$

where $h > 0$ is a control limit chosen to achieve a given ARL_0 value. An analogous charting statistic can be defined for detecting a downward mean shift.

By using the above conventional CUSUM charts, we can only know whether the process is IC or not at a given time point. It is usually unknown to us the likelihood of a potential shift, even after a signal of shift is given. To overcome this limitation, we can use the p -value of its charting statistic in its design, which is described below in the case for detecting an upward mean shift. The other cases for detecting a downward or arbitrary shift can be discussed similarly. Let C_n^{+*} be the observed value of the charting statistic C_n^+ . Then, the p -value at the n -th time point is defined by

$$P_{C_n^{+*}} = P(C_n^+ > C_n^{+*}). \quad (1)$$

As in the context of hypothesis testing, at a pre-specified significance level α , if

$$P_{C_n^{+*}} < \alpha, \quad (2)$$

then we conclude that the process is OC at the n -th time point. Otherwise, it is IC. There are several benefits to use the p -value to make decisions about the process performance. First, the control limit of the above chart is always α which has a probability interpretation. Therefore, the chart (1)-(2) is easy to construct and interpret. Second, the p -value provides a numerical measure of the likelihood of a potential shift, based on which we can take appropriate subsequent actions. For instance, if $P_{C_n^{+*}}$ is much larger than α , then we can adjust the next sample accordingly, by either waiting longer than usual for the next sample or collecting less observations at the next regular sampling time. In cases when $P_{C_n^{+*}} < \alpha$, the control chart

gives a signal of shift and the process monitoring should be stopped. In practice, however, even in such cases, the p -value $P_{C_n^{+*}}$ is helpful in taking proper post-signal actions. For instance, if $P_{C_n^{+*}}$ is much smaller than α , then the production process should be stopped immediately. In cases when $P_{C_n^{+*}}$ is only marginally smaller than α , we may want to keep the production process running, collect one more sample in a shorter-than-usual period of time (which can also be determined by our proposed dynamic sampling scheme introduced in Section 2.2), and make an appropriate decision based on the information in the new sample. Further study of the use of the p -value as a diagnostic aid will be considered in the future.

To calculate the p -value $P_{C_n^{+*}}$, we need to specify the IC distribution of C_n^+ . In cases when the IC process distribution is known, the IC distribution of C_n^+ can be determined by a Monte Carlo simulation. It is well known that this IC distribution is stable when n is large (e.g., $n \geq 50$), which is the so-called steady-state distribution in the literature (e.g., Hawkins and Olwell 1998, page 61). Therefore, in such cases, we only need to determine the IC distributions of C_n^+ for cases when n is small, and determine the steady-state distribution for approximating the IC distributions of C_n^+ in cases when n is large. These distributions can be easily tabulated by a simulation beforehand, and the p -value $P_{C_n^{+*}}$ can be computed accordingly. If it is reasonable to assume that the shift can only occur after $n \geq 50$, then only the steady-state distribution of C_n^+ is needed to compute the p -values $P_{C_n^{+*}}$, for $n \geq 50$. In cases when the IC process distribution is unknown but an IC dataset is available, the IC distribution of C_n^+ can be determined by a bootstrap procedure described as follows, which is similar to the method discussed in Chatterjee and Qiu (2009). By the bootstrap method (Efron 1979, Efron and Tibshirani 1993), we repeatedly draw observations with replacement from the IC data, which are called the resampled data, and the resampled data are used as Phase II observations to compute the values of C_n^+ , for any given n . Then, this process is repeated B times, and the B values of C_n^+ are used for estimating the IC distribution of C_n^+ , where $B > 0$ is the so-called bootstrap sample size. Again, we only need to estimate the IC distributions of C_n^+ when n is small. When n is large, we can use a single steady-state distribution to approximate all IC distributions of C_n^+ , as discussed above.

2.2 CUSUM charts using a dynamic sampling scheme

As mentioned in Section 2.1, with the CUSUM chart (1)-(2), it is natural to adjust the next sampling time according to the observed p -value $P_{C_n^{+*}}$ at the current time point n . If $P_{C_n^{+*}}$ is larger, then we can wait longer to collect the next observation. Otherwise, the next observation time should be sooner than usual. See Table 2 in Section 4 for a demonstration. Therefore, the sampling interval function $d(\cdot)$ should be an increasing function of $P_{C_n^{+*}}$. To specify this function, it is important to strike a balance between a rigid parametric form and an overly complex function. A complex function will have too many parameters to choose, while a rigid parametric form may not be sufficiently flexible to give the desired chart performance. To trade off these two considerations, we suggest choosing $d(\cdot)$ from the Box-Cox transformation family (Box and Cox 1964), because this family is flexible enough for most applications and it contains many commonly used parametric monotone functions, such as the log, linear, and quadratic functions. Some other types of transformation family, such as the Johnson transformation family, are not considered here due to their complexity and inferior performance (Qiu and Li 2011a). Define

$$d(P_{C_n^{+*}}) = a^* + b^* B_\lambda(P_{C_n^{+*}}), \quad (3)$$

where a^* and b^* are two coefficients, $B_\lambda(P_{C_n^{+*}})$ is the Box-Cox transformation of $P_{C_n^{+*}}$ defined by

$$B_\lambda(P_{C_n^{+*}}) = \begin{cases} \frac{P_{C_n^{+*}}^\lambda - 1}{\lambda} & \text{if } \lambda \neq 0 \\ \log(P_{C_n^{+*}}) & \text{if } \lambda = 0, \end{cases}$$

and $\lambda \geq 0$ is a parameter. Substituting $B_\lambda(P_{C_n^{+*}})$ into equation (3), we have

$$d(P_{C_n^{+*}}) = \begin{cases} a + b P_{C_n^{+*}}^\lambda & \text{if } \lambda \neq 0 \\ a + b \log(P_{C_n^{+*}}) & \text{if } \lambda = 0, \end{cases} \quad (4)$$

where

$$\begin{cases} a = a^* - \frac{b^*}{\lambda}, b = \frac{b^*}{\lambda} & \text{if } \lambda \neq 0 \\ a = a^*, b = b^* & \text{if } \lambda = 0. \end{cases}$$

In cases when the sampling interval remains constant in a process, the average run length (ARL) is traditionally employed in the SPC literature as a performance measure of a control chart. However, when the sampling interval is variable, the time to signal is not a constant multiple of the ARL, and thus ARL may not be appropriate for evaluating the effectiveness of

a VSI control chart. In such cases, two widely used performance measures are the average time to signal (ATS), defined as the expected value of the time interval from the start of the Phase II process monitoring to the time when a chart gives an OC signal, and the adjusted average time to signal (AATS), defined as the expected value of the time interval from the occurrence of a shift to the time when the chart gives an OC signal. When the process is IC, the ATS provides a measure of the false alarm rate (denoted ATS_0) for a control chart: the chart with a larger IC ATS would have a lower false alarm rate. When the process is OC, the AATS can be used to measure the OC performance of a chart (denoted $AATS_1$): the chart with a smaller OC AATS would perform better for detecting the related shift.

In the sampling interval function $d(\cdot)$ defined in (4), there are three parameters a , b and λ to determine. For a VSI chart, it is often required that $ATS_0 = ARL_0$. Therefore, as long as a and λ are determined, b can be determined accordingly to satisfy this requirement. For this reason, next, we discuss the selection of a and λ only.

To investigate the effect of a , let us consider cases when the IC process distribution is $N(0, 1)$, $ATS_0 = 200$, $k = 0.25$, λ values are chosen from 0, 0.5, 1, 1.5, 2, 2.5, 3, 6, and 10, and a values are chosen from 0, 0.2, 0.4, 0.6, 0.8 and 1.0. When the process mean has a shift at the initial observation time with the size 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1.0, 1.5, or 2.0, the $AATS_1$ values of the chart (2)-(3) with $\lambda = 0$ and 0.5, computed from 10,000 replicated simulation runs, are presented in Figure 1. The $AATS_1$ values of the chart (2)-(3) with λ values chosen from 0, 0.5, 1, 1.5, 2, 2.5, 3, 6, and 10, are presented in Figure S.1 of the online supplementary material. Note that in the VSR problems such as the one discussed here, the basic time unit is assumed to be well defined beforehand, and we collect one sample every one basic time unit under a FSR scheme. Therefore, by assuming $ATS_0 = 200$, we actually assume that the IC ATS is 200 basic time units. In this example, the value of a is chosen in the interval $[0, 1]$ for the following reason. If a is chosen a negative number, then the sampling interval by (4) could be negative. If a is chosen larger than 1, then the sampling interval would be consistently larger than 1 in cases when $\lambda > 0$, which is not reasonable because the chart (2)-(3) is supposed to take the next sample sooner than usual when the observed p -value $P_{C_n^{+*}}$ is close to α . From Figure 1 and Figure S.1, it can be seen that (i) when $\lambda = 0$, the $AATS_1$ value decreases when a increases, and (ii) when $\lambda > 0$, the $AATS_1$ value increases when a increases. Based on this study, we should choose $a = 1$ when $\lambda = 0$, and $a = 0$ when $\lambda > 0$.

To investigate the effect of λ , we consider nine λ values $\{0, 0.5, 1, 1.5, 2, 2.5, 3, 6, 10\}$ in

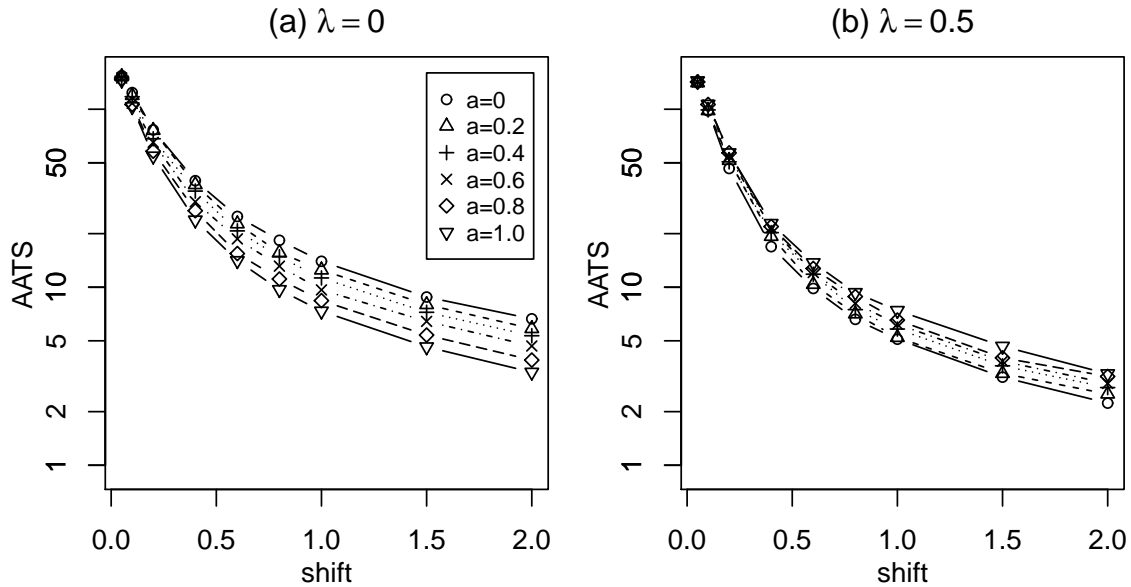


Figure 1: $AATS_1$ values of the chart (2)–(3) when the IC process distribution is $N(0, 1)$, the mean shift size at the initial observation time changes from 0.05 to 2.0, $ATS_0 = 200$, and $k = 0.25$.

the interval $[0, 10]$. By the results in the above example, a is chosen 1 when $\lambda = 0$, and 0 when $\lambda > 0$. All other setup is the same as that in the above example. The computed $AATS_1$ values of the chart (2)–(3) are shown in Figure 2(a). From the plot, it seems that the performance of the chart gets better when λ increases; but, the performance does not change much when $\lambda \geq 2$. To further investigate this issue, we consider another nine λ values $\{1.6, 1.7, 1.8, 1.9, 2, 2.1, 2.2, 2.3, 2.4\}$ around 2. The corresponding $AATS_1$ values are shown in Figure 2(b), from which it can be seen that $AATS_1$ is indeed stable when $\lambda \geq 2$. Based on these results and on the practical guideline to choose a simple number for λ when using the Box-Cox transformation (e.g., Weisberg 2005, Section 7.1), it seems that $\lambda = 2$ is a reasonable choice.

Based on results in Figures 1 and 2, in cases when the IC process distribution is $N(0, 1)$, we suggest choosing the sampling interval function of the chart (2)–(3) to be

$$d(P_{C_n^+}) = b \cdot P_{C_n^+}^2. \quad (5)$$

Note that this sampling interval function is based on our empirical study. Much more future research is required to further justify this choice using mathematically more rigorous arguments. It should be pointed out that, there is only one parameter b in equation (5), which can be determined to satisfy the requirement that $ATS_0 = ARL_0$. As a comparison, in a conventional VSI control chart (reviewed in Section 3 below), there are usually two or more parameters involved in its sampling interval function. Therefore, our proposed control chart described in

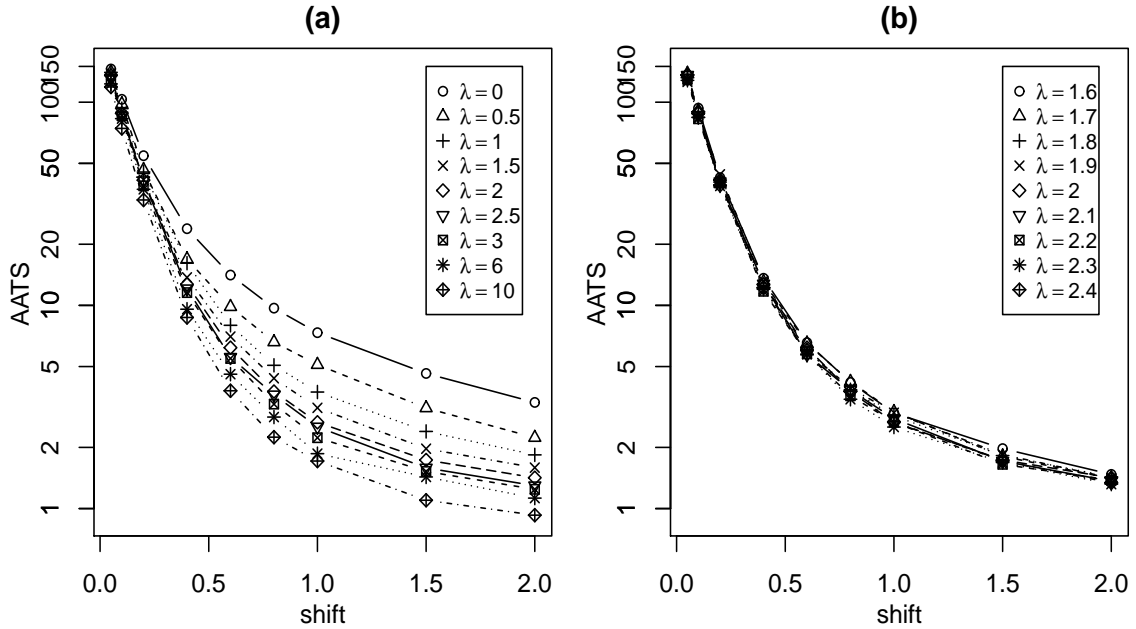


Figure 2: $AATS_1$ values of the chart (2)–(3) when the IC process distribution is $N(0, 1)$, the mean shift size at the initial observation time changes from 0.05 to 2.0, $ATS_0 = 200$, and $k = 0.25$. Plot (b) considers a narrower range of λ values.

(1)–(5) would be easier to use, compared to the conventional VSI control chart. Also, although theoretically speaking, the sampling interval function defined in (5) can take any values in $(0, b)$, in practice it can only take integer multiples of the smallest time unit (e.g., one hour) in a specific application. Therefore, the computed sampling intervals need to be rounded when necessary.

From its definition, our charting statistic $P_{C_n^+}$ is uniquely determined by the distribution of the conventional CUSUM charting statistic C_n^+ , and the latter converges to a steady-state distribution when n is reasonably large. Therefore, we expect that the IC and OC performance of the chart (2)–(3) would also have the steady-state property. To demonstrate this, in the example of Figure 2, we consider cases when $a = 0$, $\lambda = 0.5, 1, 1.5$, or 2 , the shift time $\tau = 1, 5, 10$, or 50 , and the other setup is unchanged. The computed AATS values of the chart (2)–(3) with $\lambda = 2$ are shown in Figure 3 and the corresponding results with $\lambda = 0.5, 1$, and 1.5 are shown in Figure S.2 of the online supplementary file. From these figures, it can be seen that (i) τ does have a substantial impact on the OC performance of the chart when it is small, and (ii) the OC performance of the chart is stable when $\tau \geq 10$.




Figure 3: $AATS_1$ values of the chart (2)–(3) when the IC process distribution is $N(0, 1)$, the mean shift size at the initial observation time changes from 0.05 to 2.0, $ATS_0 = 200$, $k = 0.25$, the shift time $\tau = 1, 5, 10$, or 50, $a = 0$, and $\lambda = 2$.

2.3 Adaptive selection of the reference value

To use our proposed control chart (2)–(5), we still need to choose the reference value k that is contained in the definition of the charting statistic C_n^+ . For the conventional CUSUM chart with a constant sampling interval function, practitioners often pre-specify the value of k . However, the resulting CUSUM chart is only optimal for detecting a specific shift of the size $2k$, and it may not perform well for detecting shifts of other sizes. To overcome this limitation, Sparks (2000) suggested estimating the shift size at each time point and then choosing the reference value k accordingly. This scheme for choosing k is often called the adaptive selection scheme in the literature. In this paper, we suggest using the adaptive selection scheme in our proposed control chart (2)–(5), described briefly as follows. Let

$$\hat{\delta}_n = \max \left\{ \delta_{\min}, (1 - r)\hat{\delta}_{n-1} + rX_n \right\}$$

be an estimator of the size of a potential mean shift at the current time point, where $\delta_{\min} > 0$ is the minimum shift size that we are interested in detecting, $\hat{\delta}_0 = \delta_{\min}$, and $0 < r \leq 1$ is a

weighting parameter. Then, we define $k_n = \widehat{\delta}_n/2$, and the resulting charting statistic becomes

$$\begin{cases} C_0^+ = 0, \\ C_n^+ = \max\left(0, C_{n-1}^+ + (X_n - \mu_0 - k_n)/h_n\right), \end{cases} \quad (6)$$

where $h_n > 0$ is a control limit. Shu and Jiang (2006) provided the following formula to compute the corresponding control limit h_n such that a pre-specified ARL_0 could be approximately reached:

$$h_n = \frac{\log(1 + 2k_n^2 \cdot ARL_0 + 2.332k_n)}{2k_n} - 1.166,$$

which is a modification of the formula proposed by Siegmund (1985). Both Sparks (2000) and Shu and Jiang (2006) have demonstrated that the CUSUM chart with the above adaptive selection scheme could perform well in various cases, and they also provided some practical guidelines for choosing the parameters δ_{\min} and r .

Next, we demonstrate numerically that the charting statistic defined in (6) with the adaptively selected reference value k_n still has the steady-state property. To this end, we consider four combinations of (r, δ_{\min}) : (0.2, 0.5), (0.2, 1.0), (0.1, 0.5), (0.1, 1.0). For each combination, the empirical distribution of C_n^+ defined by (6) is obtained by a Monte Carlo simulation with 1 million replications when $n = 10, 20, 50, 100, 200$ and 500 . Then, the p -values $P_{C_n^{+*}}$ for various observed values C_n^{+*} of C_n^+ can be computed, and those when $(r, \delta_{\min}) = (0.2, 0.5)$ and $(0.2, 1.0)$ are shown in the two plots of Figure 4 here and those when $(r, \delta_{\min}) = (0.1, 0.5)$ and $(0.1, 1.0)$ are presented in Figure S.3 of the online supplementary file. (Note that $1 - P_{C_n^{+*}}$ is just the empirical distribution of C_n^+ .) From the two figures, it can be seen that (i) the empirical distribution of C_n^+ is almost identical when $n \geq 50$ in all cases considered, (ii) the empirical distribution of C_n^+ depends on n in a substantial way only in cases when n is small and when δ_{\min} is small as well, and (iii) when δ_{\min} is relatively large (i.e., $\delta_{\min} = 1.0$ in the current example), the empirical distribution of C_n^+ is almost identical when n is as small as 10. This example also shows that the weighting parameter r has little effect on the distribution of C_n^+ .

In the above example, the empirical distribution of C_n^+ , or the p -values $P_{C_n^{+*}}$, are computed using Fortran codes (available from the online supplementary material) together with the IMSL subroutine ‘‘cpsec’’. For each combination of (r, δ_{\min}) , the CPU times for finding the p -values based on 1 million replicated simulations are 51.54, 52.82, 53.95, 54.88, 56.74 and 60.82 seconds, respectively, in cases when $n = 10, 20, 50, 100, 200$ and 500 , on an Intel 2 with CPU processor




Figure 4: p -value $P_{C_n^+}$ computed from the empirical distribution of C_n^+ defined by (6) in cases when $n = 10, 20, 50, 100, 200$ and 500 .

2.4 GHz. Note that, in practice, to use our proposed method, we only need to specify the empirical distributions of C_n^+ in cases when $n \leq 50$ and these empirical distributions can be tabulated and saved in our computers beforehand. Therefore, it is actually quite convenient to use our proposed method.

As a summary, our proposed CUSUM chart uses the charting statistic defined in (6), and its sampling interval function is defined in (5), where b is chosen to satisfy the requirement $ATS_0 = ARL_0$. The chart gives a signal of process mean shift if the expression (2) is true. This chart is called the dynamic-sampling CUSUM (DyS-CUSUM) chart hereafter. An illustration of the implementation of the DyS-CUSUM chart is given by Table 2 in Section 4.

3 Simulation Study

In this section, we present some simulation results to evaluate the numerical performance of our proposed chart DyS-CUSUM. In this section, all AATS values are computed based on 10,000 replicated simulation runs. In such cases, the standard errors of the computed AATS values are usually less than 2% of the AATS estimates, enabling us to draw reasonable conclusions about the performance of the DyS-CUSUM chart. The empirical distribution of C_n^+ is computed based on a million replicated simulation runs, as in the previous section.

In our numerical study, we compare the DyS-CUSUM chart with the VSI adaptive CUSUM chart suggested by Luo et al. (2009), called the VSI-ACUSUM chart hereafter. The VSI-ACUSUM chart uses the conventional CUSUM charting statistic described in Section 2.1, the VSI sampling scheme described below, and the adaptive selection of the reference value discussed in Section 2.3. The VSI sampling scheme used by VSI-ACUSUM is the conventional 2-interval sampling scheme. Let d_1 and d_2 be two possible sampling intervals with $0 < d_1 < d_2$. Then, the corresponding sampling interval function $d(\cdot)$ is defined to be

$$d(C_n^+) = \begin{cases} d_1, & \text{if } C_n^+ \in R_w \\ d_2, & \text{if } C_n^+ \in R_c, \end{cases} \quad (7)$$

where $R_c = [0, h_1]$ is the central region, $R_w = (h_1, h_2]$ is the warning region, $0 < h_1 < h_2$ are two control limits, and they are chosen such that (i) a given value of ARL_0 is achieved, and (ii) ARL_0 is about the same as ATS_0 . Existing research on VSI (e.g., Costa 1998, Wu et al. 2007, Reynolds and Arnold 2001, Stoumbos et al. 2011) shows that d_1 should be chosen as small as possible, and d_2 should be chosen as large as possible. Reynolds et al. (1990) showed that the VSI sampling scheme with $d_1 = 0.1$ and $d_2 = 1.9$ can give good results in various cases, which was used in Luo et al. (2009) and is used here as well.

From the above description, it can be seen that the major differences between the VSI-ACUSUM chart and our proposed DyS-CUSUM chart are that (i) the VSI-ACUSUM chart uses a conventional 2-interval sampling scheme while we use a dynamic sampling scheme, and (ii) the VSI-ACUSUM chart uses the conventional control limits in its design and we use the p -value of the charting statistic in the design of the DyS-CUSUM chart.

In our numerical study, the ATS_0 values of both related charts are set to be 400, as in Luo et al. (2009). We assume that the IC process distribution is $N(0, 1)$, and the mean shift changes its value among 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1.0, 1.5, and 2.0. For both the DyS-CUSUM and VSI-ACUSUM charts, their reference values are chosen adaptively by the procedure described in Section 2.3. Following Shu and Jiang (2006), the weighting parameter r and minimum shift size δ_{min} are set to be $r = 0.2$ and $\delta_{min} = 0.05$. Then, the b value is computed to be 3.1562 so that $ATS_0 = ARL_0 = 400$. To evaluate the performance of the related control charts, besides the $AATS_1$ values, we also consider the so-called integral of the relative AATS (IRAATS) measure to evaluate the overall performance of the control charts in detecting different shifts, which was

first suggested by Luo et al. (2009). For a control chart C , its IRAATS value is defined by

$$\text{IRAATS} = \frac{1}{p} \sum_{j=1}^p \frac{\text{AATS}_{1,C}(\delta_j)}{\text{AATS}_{1,\min}(\delta_j)},$$

where $\text{AATS}_{1,C}(\delta_j)$ is the AATS_1 value of the control chart C for detecting the shift δ_j , $\text{AATS}_{1,\min}(\delta_j)$ is the smallest AATS_1 value of all control charts considered for detecting the shift δ_j , and p is the total number of shifts considered. By this measure, a control chart with a smaller IRAATS value is considered to be more effective in detecting the related shifts. Note that one may intentionally lower the IRAATS value of a control chart by including more cases with large shifts, because the chart would have small AATS_1 values for detecting large shifts. Similarly, one can intentionally increase the IRAATS value of the chart by considering more cases with small shifts.

Table 1: AATS_1 values of the control charts DyS-CUSUM and VSI-ACUSUM when they are used for detecting mean shifts δ_j . It is assumed that the IC process distribution is $N(0, 1)$ and $\text{ATS}_0 = 400$ for both control charts.

δ_j	DyS-CUSUM		VSI-ACUSUM	
	adaptive	k=0.2	adaptive	k=0.2
0.05	198.26	244.45	192.47	268.95
0.10	97.48	137.74	97.79	164.17
0.20	33.84	49.96	36.29	56.07
0.40	12.76	13.25	12.32	12.10
0.60	6.49	6.40	7.10	6.00
0.80	4.27	3.79	4.94	3.82
1.00	2.70	2.84	3.61	2.88
1.50	1.69	1.82	2.09	1.86
2.00	1.35	1.42	1.50	1.47
<i>IRAATS</i>	1.03	1.17	1.14	1.22

The results of the two charts are presented in Table 1 in columns labeled by “adaptive.” It can be seen from the table that DyS-CUSUM is better than VSI-ACUSUM in all cases except the one with $\delta_j = 0.05$. The IRAATS value of the DyS-CUSUM chart is about 10.5% smaller than that of the VSI-ACUSUM chart. Therefore, the overall performance of the DyS-CUSUM chart is better than that of the VSI-ACUSUM chart. As mentioned above, the main difference between the charts DyS-CUSUM and VSI-ACUSUM is that the former uses the dynamic sampling scheme while the latter uses the regular 2-interval sampling scheme. So, this example demonstrates that the dynamic sampling scheme has certain advantage, compared to the regular 2-interval sampling scheme. To further demonstrate this result, we also present the corresponding results of the two charts when their reference values are both fixed at $k = 0.2$ in the same table. It can be seen that similar conclusions can be made in such cases.

In the next example, we investigate the effectiveness of the dynamic sampling scheme used in the DyS-CUSUM chart. To this end, in the control chart (1)-(2) using the p -value of its charting statistic, we consider three different sampling schemes: (i) $d(C_n^+) = 1$, (ii) $d(C_n^+)$ is the conventional 2-interval scheme defined by equation (7) with $d_1 = 0.1$ and $d_2 = 1.9$ as in the VSI-ACUSUM chart, and (iii) the dynamic sampling scheme defined in equation (5). In each case, the reference value k is chosen by the adaptive scheme, and the other parameters are chosen such that $ATS_0 = 400$. In this study, we assume that the IC process distribution is unknown. Instead, we have 2,000 IC observations from the $N(0, 1)$ distribution. Therefore, in the control chart (1)-(2), the bootstrap approach with $B = 1,000,000$ is used for computing the p -value, as described in the last paragraph of Section 2.1. The $AATS_1$ values of the chart in various different cases are presented in Figure 5(a). From the plot, it can be seen that the conventional 2-interval VSI scheme does improve the constant sampling scheme, but the dynamic sampling scheme can further improve the conventional 2-interval VSI scheme.

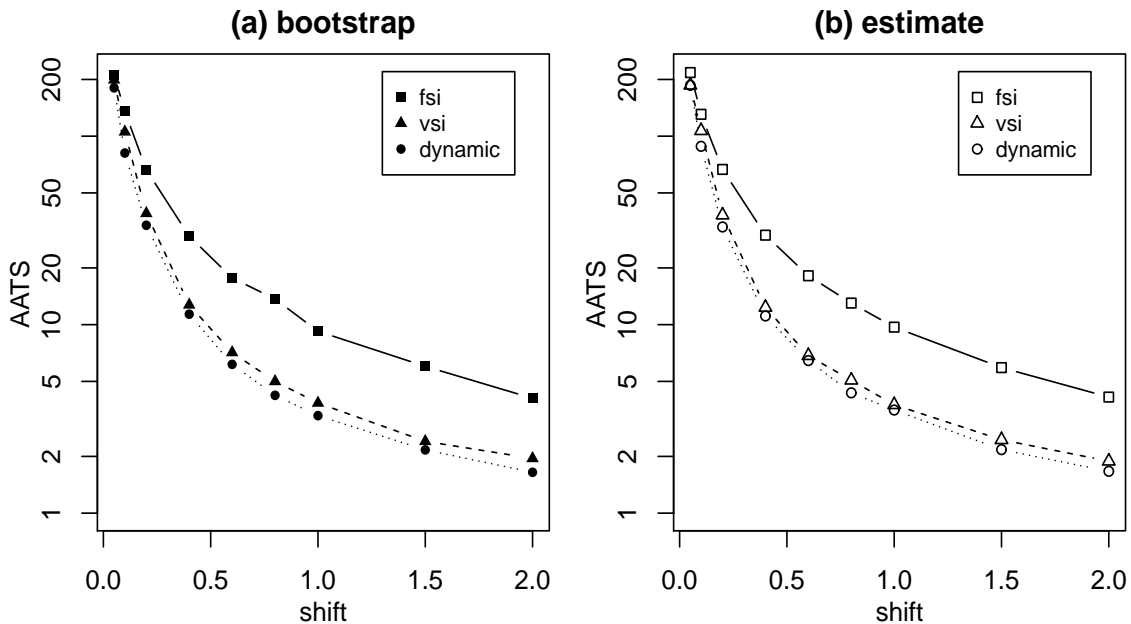


Figure 5: $AATS_1$ values of the control chart (1)–(2) when its reference value k is chosen by the adaptive selection scheme, and the sampling interval function $d(\cdot)$ is chosen by (i) the fixed sampling interval (FSI) scheme $d(C_n^+) = 1$, (ii) the conventional 2-interval VSI scheme defined by equation (7), and (iii) the dynamic VSI scheme defined by equation (5). In plot (a), the p -value of the charting statistic is computed by the bootstrap approach with $B = 1,000,000$. In plot (b), the p -value is computed by the distribution estimation approach. In each case, ATS_0 is fixed at 400.

As an alternative to the bootstrap approach, one can also estimate the IC distribution parameters from an IC dataset in cases when the IC distribution family is given. Then, the p -value

of the charting statistic can be computed as if the IC distribution is known. To investigate this approach, we assume that the IC distribution is univariate normal, and its mean and standard deviation are estimated from 2,000 IC observations generated from the $N(0, 1)$ distribution. Then, we treat the estimated mean and standard deviation as their true values and compute the p -value of the charting statistic based on this estimated IC distribution. The $AATS_1$ values of the chart with estimated parameters in various different cases are presented in Figure 5(b). From the plot, we can see that the same conclusions as those from Figure 5(a) can be made here. To compare the bootstrap approach with this distribution estimation approach, we show the $AATS_1$ values by both approaches in cases with the constant sampling scheme, the 2-interval VSI scheme and the dynamic sampling scheme in Figure S.4(a)-(c) of the online supplementary file, respectively. From that figure, it can be seen that our bootstrap approach has nearly the same accuracy as the distribution estimation approach. But, the bootstrap approach does not require to specify the IC distribution family, while the distribution estimation approach does.

In our proposed DyS-CUSUM chart, the process observations are assumed to be independent. In practice, however, this assumption is often invalid, especially when the sampling times are close to each other. In the next example, we investigate the performance of the DyS-CUSUM chart when process observations are correlated. More specifically, assume that the current observation X_n and the previous observation X_{n-1} are t time units apart, and they follow the following time series model discussed in Reynolds et al. (1996):

$$X_n = (1 - \phi^t)\mu_0 + \phi^t X_{n-1} + \epsilon_n, \quad n = 1, 2, \dots, \quad (8)$$

where μ_0 is the IC process mean, ϵ_n is the standard normal random error, and $0 < \phi < 1$ is an autoregressive parameter. From this model, it is easy to check that the correlation coefficient between X_{n-1} and X_n is ϕ^t . It approaches 0 when the time difference t gets large, and approaches 1 when t tends to 0. In cases when $ATS_0 = 200$, $\phi = 0, 0.25, 0.5$, and 0.75 , and the mean shift changes its value among 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1.0, 1.5, and 2.0, and the other setup is the same as those in Figure 1, the AATS values of DyS-CUSUM are shown in Figure 6. From the figure, we can make three conclusions. First, as the correlation coefficient ϕ becomes larger, the actual ATS_0 value when the process is IC is farther away from the nominal ATS_0 value of 200. For instance, when $\phi = 0.75$, the actual ATS_0 value is 176, which is substantially below 200. Second, when ϕ increases, the actual $AATS_1$ value would increase as well when detecting most mean shifts. Third, when $\phi \leq 0.5$, it seems that the impact of the autocorrelation on the performance of the DyS-CUSUM chart is limited.




Figure 6: $AATS_1$ values of the DyS-CUSUM chart when process observations follow the time series model (8) with various values of the autocorrelation parameter ϕ , and the mean shift changes its value from 0.05 to 2.0.

4 Real Data Example

In this section, we demonstrate our proposed DyS-CUSUM chart using a real-data obtained from a chemical process. The data set contains 149 readings of triglyceride content of chemical products, which is described in more detail in Chapter 3 of Hawkins and Olwell (1998). The data can be downloaded from <http://www.stat.umn.edu/cusum/data.htm>.

Hawkins and Olwell (1998) have shown that the process mean appears to be well within its allowable range in the early part of the data set. Based on that result, we use the first 75 observations as an IC dataset, and the remaining observations are used for testing. The Shapiro-Wilk test for normality, which can be accomplished by the **R** function `shapiro.test`, gives a p -value of 0.03852 when it is applied to the first 75 observations. This test shows that the IC process distribution is marginally significantly different from a normal distribution. Therefore, in this example, we use both the bootstrap approach, which does not require the specification of a parametric form for the IC process distribution, and the distribution estimation approach, which assumes that the IC process distribution is a normal distribution, when computing the p -value of the charting statistic C_n^+ , as discussed in the example of Figure 5. The two versions of the DyS-CUSUM chart are denoted as DyS-CUSUM-B and DyS-CUSUM-E, respectively. In both versions, we choose $ARL_0 = 400$, $r = 0.2$ and $\delta_{min} = 0.05$, as in Section 3. The significance

level α used in (2) is computed to be 0.025 and 0.024, respectively, for DyS-CUSUM-B and DyS-CUSUM-E. Then, the DyS-CUSUM-B chart is shown in Figure 7 and the DyS-CUSUM-E chart is shown in Figure S.5 of the online supplementary file. In both charts, the horizontal dashed lines denote the corresponding significance levels. From the two figures, we can see that both versions of the DyS-CUSUM chart give signals of process mean shift at the 123rd time point and the subsequent p -values are all well below the significance levels. Therefore, these signals are convincing enough, and they are also consistent with the findings in Hawkins and Olwell (1998).

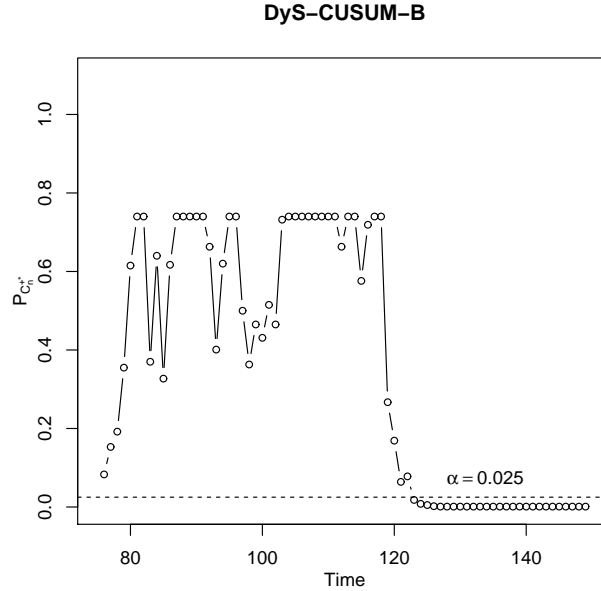


Figure 7: Control chart DyS-CUSUM-B for monitoring a univariate chemical process. The horizontal dashed line in the plot denotes the significance level of the chart such that $ARL_0 = 400$.

At the end of this section, we would like to use this example to illustrate the implementation of the DyS-CUSUM chart. In the illustration, we use the bootstrap approach only when computing the p -value of the charting statistic C_n^+ . Note that, in the original data, observations are collected at equally spaced time points. However, for the illustration purpose, we assume that they are collected using a dynamic sampling scheme specified by the sampling interval function (5) of our DyS-CUSUM chart. The coefficient b in equation (5) is computed to be 3.3711 such that $ATS_0 = ARL_0 = 400$. For the testing data (i.e., 76th to 149th observations), the adaptively selected reference values k_n , the charting statistic values C_n^{+*} , the corresponding p -values $P_{C_n^{+*}}$, and the sampling intervals $d(P_{C_n^{+*}})$ of the DyS-CUSUM chart are presented in Table 2. Note that, such values are not given for observations 123-149, because an OC signal

is already given at $n = 123$. From Table 2, after the 76th observation is obtained, the value of $P_{C_n^{+*}}$ is computed to be 0.083, which is quite close to the significance level. Therefore, we collect the 77th observation after $b \cdot P_{C_n^{+*}}^2 = 0.0235$ time units, which is much sooner than the regular sampling interval of 1 time unit. At the time point $n = 81$, the p -value $P_{C_n^{+*}}$ is computed to be 0.740, which is quite large, indicating that a mean shift is unlikely at that time point. The computed value of $d(P_{C_n^{+*}})$ is 1.8435, indicating that we collect the next observation after 1.8435 time units. The results at other time points can be explained in a similar way.

Table 2: Values of k_n , C_n^{+*} , $P_{C_n^{+*}}$ and $d(P_{C_n^{+*}})$ of the DyS-CUSUM chart based on the bootstrap approach when computing the p -value of the charting statistic C_n^+ .

n	k_n	C_n^{+*}	$P_{C_n^{+*}}$	$d(P_{C_n^{+*}})$	n	k_n	C_n^{+*}	$P_{C_n^{+*}}$	$d(P_{C_n^{+*}})$
76	0.466	0.9589	0.083	0.0235	113	0.250	0.0000	0.740	1.8435
77	0.250	0.7034	0.153	0.0794	114	0.250	0.0000	0.740	1.8435
78	0.250	0.5919	0.192	0.1236	115	0.279	0.0816	0.576	1.1203
79	0.250	0.2884	0.355	0.4260	116	0.250	0.0181	0.719	1.7402
80	0.250	0.0809	0.615	1.2729	117	0.250	0.0000	0.740	1.8435
81	0.250	0.0000	0.740	1.8435	118	0.250	0.0000	0.740	1.8435
82	0.250	0.0000	0.740	1.8435	119	0.441	0.4295	0.267	0.2412
83	0.376	0.2690	0.370	0.4627	120	0.496	0.6552	0.169	0.0968
84	0.250	0.0615	0.640	1.3829	121	0.605	1.0732	0.064	0.0135
85	0.376	0.3306	0.327	0.3615	122	0.499	0.9883	0.078	0.0202
86	0.250	0.0750	0.617	1.2854	123	0.672	1.6264	0.018	*
87	0.250	0.0000	0.740	1.8435	124	0.713	1.9695	0.008	*
88	0.250	0.0000	0.740	1.8435	125	0.682	2.1055	0.005	*
89	0.250	0.0000	0.740	1.8435	126	0.754	2.5642	0.002	*
90	0.250	0.0000	0.740	1.8435	127	0.876	3.3047	0.001	*
91	0.250	0.0000	0.740	1.8435	128	0.845	3.5331	0.001	*
92	0.250	0.0326	0.663	1.4840	129	0.916	4.1563	0.001	*
93	0.344	0.2322	0.401	0.5434	130	0.941	4.6476	0.001	*
94	0.250	0.0727	0.620	1.2937	131	0.929	5.0004	0.001	*
95	0.250	0.0000	0.740	1.8435	132	0.984	5.6416	0.001	*
96	0.250	0.0000	0.740	1.8435	133	0.931	5.8570	0.001	*
97	0.311	0.1371	0.500	0.8444	134	0.953	6.3496	0.001	*
98	0.361	0.2794	0.363	0.4429	135	0.680	5.8784	0.001	*
99	0.250	0.1679	0.465	0.7304	136	0.752	6.3367	0.001	*
100	0.250	0.2005	0.431	0.6247	137	0.649	6.2826	0.001	*
101	0.250	0.1371	0.515	0.8923	138	0.824	7.1202	0.001	*
102	0.250	0.1697	0.465	0.7273	139	1.126	8.9729	0.001	*
103	0.250	0.0102	0.732	1.8038	140	0.980	8.8884	0.001	*
104	0.250	0.0000	0.740	1.8435	141	1.347	11.6756	0.001	*
105	0.250	0.0000	0.740	1.8435	142	1.124	11.3328	0.001	*
106	0.250	0.0000	0.740	1.8435	143	1.237	12.5827	0.001	*
107	0.250	0.0000	0.740	1.8435	144	0.940	11.9645	0.001	*
108	0.250	0.0000	0.740	1.8435	145	1.089	13.1164	0.001	*
109	0.250	0.0000	0.740	1.8435	146	1.015	13.3133	0.001	*
110	0.250	0.0000	0.740	1.8435	147	1.117	14.3168	0.001	*
111	0.250	0.0000	0.740	1.8435	148	1.134	14.9881	0.001	*
112	0.250	0.0326	0.663	1.4840	149	1.051	15.1752	0.001	*

5 Concluding Remarks

In this paper, we have presented a CUSUM chart designed by the p -value of its charting statistic using a dynamic sampling scheme. When using this chart, users can get a measure of the likelihood of a potential shift so that a subsequent action can be taken accordingly. One natural subsequent action is the dynamic sampling scheme that is described in the paper, by which the sampling interval for collecting the next observation depends on the p -value computed at the current time point. If the p -value is larger, then the sampling interval would be longer. In this paper, we have demonstrated that the dynamic sampling scheme can improve the conventional 2-interval variable sampling scheme that is extensively studied in the literature.

Although our proposed method is demonstrated using the univariate CUSUM chart (1)-(2) for detecting a process mean shift in this paper, it is actually quite general and can be applied to other control charts, such as the univariate and multivariate Shewhart and EWMA charts. It can also be applied to control charts for detecting other types of process distributional shifts, such as the process variance shifts. For instance, if it is confirmed that the process distribution is non-normal, then certain nonparametric control charts can be considered (cf., Qiu and Hawkins 2001, Qiu and Li 2011b), and the p -value and dynamic sampling scheme can be applied to such nonparametric control charts accordingly.

We should point out that the proposed CUSUM chart with dynamic sampling intervals has some disadvantages. First, the sampling intervals specified by the dynamic sampling interval function in equation (5) change continuously, and it may bring some inconvenience to real-world implementation. Second, the implementation of the proposed control chart may require the use of a computer for computing the p -value of the charting statistic at each time point. This disadvantage might be minor nowadays because many manufacturing facilities have been equipped with computers. Third, our scheme is based on the usual assumption that successive measurements are independent, which, however, may become less plausible as the sampling frequency increases, since measurements made at closer time points are more likely to be correlated. It requires much future research to handle correlated data when using our proposed CUSUM chart.

Supplementary Materials

supplemental.pdf: This pdf file presents some extra numerical results.

ComputerCodesAndData.zip: This zip file contains Fortran source codes of our proposed method and the triglyceride data used in the paper.

Acknowledgments

The authors are grateful to the editor, the associate editor and two anonymous referees for their valuable comments that have greatly improved the paper. Part of this research is finished during Li's visit to School of Statistics at The University of Minnesota, whose hospitality is appreciated. The research is supported in part by an NSF grant, by the Natural Sciences Foundation of China grants 11201246, 11071128 and 11131002, the RFDP of China Grant 20110031110002, and the Office of International Programs at The University of Minnesota.

References

- [1] Benjamini, Y., and Kling, Y. (1999), "A Look at Statistical Process Control through the P-Value," *Technical report RPCSOR-99-08*, Tel Aviv University, Israel.
- [2] Benjamini, Y., and Kling, Y. (2002), "Aspects of Multiplicity in Statistical Process Control," Available at <http://businessken.co.uk/Documents/Aspects%20of%20Multiplicity%20in%20SPC.pdf>.
- [3] Benjamini, Y., and Kling, Y. (2007), "The P-Valued Chart-A Unified Approach to Statistical Process Control Chart Presentation," Available at <http://businessken.co.uk/Documents/SPCpvalue.pdf>.
- [4] Box, G. E. P., and Cox, D. R. (1964), "An Analysis of Transformations," *Journal of the Royal Statistical Society Series B*, 26 (2), 211-252.
- [5] Brook, D., and Evans, D. A. (1972), "An Approach to the Probability Distribution of Cusum Run Length," *Biometrika*, 59, 539-549.
- [6] Chatterjee, S., and Qiu, P. (2009), "Distribution-Free Cumulative Sum Control Charts Using Bootstrap-Based Control Limits," *The Annals of Applied Statistics*, 3 (1), 349-369.
- [7] Chen, G., Cheng, S., and Xie, H. (2001), "Monitoring Process Mean and Variability with One EWMA Chart," *Journal of Quality Technology*, 33, 223-233.
- [8] Chen, G., and Zhang, L. (2005), "An Extended EWMA Mean Chart," *Quality Technology and Quantitative Management*, 2, 39-52.
- [9] Costa, A. F. B., (1998), "Joint \bar{X} and R Charts with Variable Parameters," *IIE Transactions*, 30, 505-514.
- [10] Efron, B., (1979), "Bootstrap Methods: Another Look at the Jackknife," *The Annals of Statistics*, 7, 1-26.
- [11] Efron, B., and Tibshirani, R. J. (1993), *An Introduction to the Bootstrap*, Chapman and Hall: Boca Raton, FL.
- [12] Grigg, O. A., and Spiegelhalter, D. J., (2008), "An Empirical Approximation to the Null Unbounded Steady-State Distribution of the Cumulative Sum Statistic," *Technometrics*, 50, 501-511.

- [13] Hawkins, D. M., and Olwell, D. H. (1998), *Cumulative Sum Charts and Charting for Quality Improvement*, Berlin: Springer.
- [14] Li, Y., and Tsung, F., (2009), "False Discovery Rate-Adjusted Charting Schemes for Multistage Process Monitoring and Fault Identification," *Technometrics*, 51, 186-205.
- [15] Li, Z., Qiu, P., Chatterjee, S., and Wang, Z. (2013), "Using P-Values To Design Statistical Process Control Charts," *Statistical Papers*, 54, 523-539.
- [16] Lorden, G., (1971), "Procedures for Reacting to a Change in Distribution," *The Annals of Mathematical Statistics*, 42, 1897-1908.
- [17] Lucas, J. M., and Saccucci, M. S., (1990), "Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements," *Technometrics*, 32, 1-12.
- [18] Luo, Y., Li, Z., and Wang, Z. (2009), "Adaptive CUSUM Control Chart with Variable Sampling Intervals," *Computational Statistics and Data Analysis*, 53, 2693-2701.
- [19] Montgomery, D. C. (2007), "SPC Research—Current Trends," *Quality and Reliability Engineering International*, 23, 515-516.
- [20] Montgomery, D. C. (2009), *Introduction to Statistical Quality Control (6th ed.)*, New York: John Wiley & Sons.
- [21] Moustakides, G. V. (1986), "Optimal Stopping Time for Detecting Changes in Distributions," *The Annals of Statistics*, 14, 1379-1387.
- [22] Page, E. S. (1954), "Continuous Inspection Schemes," *Biometrika*, 42, 243-254.
- [23] Pollack, M. (1985), "Optimal Detection of A Change in Distribution," *The Annals of Statistics*, 13, 206-227.
- [24] Qiu, P. (2013), *Introduction to Statistical Process Control*, London: Chapman & Hall/CRC.
- [25] Qiu, P., and Hawkins, D. (2001), "A Rank Based Multivariate CUSUM Procedure," *Technometrics*, 43, 120-132.
- [26] Qiu, P., and Li, Z. (2011a), "Distribution-Free Monitoring of Univariate Processes," *Statistics and Probability Letters*, 81, 1833-1840.
- [27] Qiu, P., and Li, Z. (2011b), "On Nonparametric Statistical Process Control Of Univariate Processes," *Technometrics*, 53, 390-405.
- [28] Reynolds, M. R., Jr., Amin, R. W., and Arnold, J. C. (1990), "CUSUM Charts with Variable Sampling Intervals," *Technometrics*, 32 (4), 371-384.
- [29] Reynolds, M. R., Jr., and Arnold, J. C. (2001), "EWMA Control Charts with Variable Sample Sizes and Variable Sampling Intervals," *IIE Transactions*, 33, 511-530.
- [30] Reynolds, M. R., Jr., Arnold, J. C., and Baik, J. W., (1996), "Variable Sampling Interval \bar{X} Charts in the Presence of Correlation," *Journal of Quality Technology*, 28, 12-30.
- [31] Roberts, G. C. (1959), "Control Chart Tests Based on Geometric Moving Average," *Technometrics*, 1, 239-250.
- [32] Shewhart, W. A. (1931), *Economic Control of Quality of Manufactured Product*, New York: Van Nostrand.
- [33] Shu, L. J., and Jiang, W. (2006), "A Markov Chain Model for the Adaptive CUSUM Control Chart," *Journal of Quality Technology*, 38, 135-147.
- [34] Siegmund, D. (1985), *Sequential Analysis: Tests and Confidence Intervals*, New York: Springer-Verlag.
- [35] Sparks, R. S., (2000), "CUSUM Charts for Signalling Varying Locations Shifts," *Journal of Quality Technology*, 32, 157-171.
- [36] Stoumbos, Z. G., Mittenthal, J., and Runger, G. C., (2011), "Steady-state-optimal Adaptive Control Charts Based on Variable Sampling Intervals," *Stochastic Analysis and Applications*, 19 (6), 1025-1057.
- [37] Weisberg, S. (2005), *Applied Linear Regression (3rd edition)*, New York: John Wiley & Sons.

- [38] Woodall, W. H. (2000), "Controversies and Contradictions in Statistical Process Control," *Journal of Quality Technology*, 32, 341-350.
- [39] Wu, Z., Zhang, S., and Wang, P. (2007), "A CUSUM Scheme with Variable Sample Sizes and Sampling Intervals for Monitoring the Process Mean and Variance," *Quality and Reliability Engineering International*, 23, 157-170.
- [40] Yeh, A. B., Lin, D. K.-J., and Venkataramani, C. (2004), "Unified CUSUM Control Charts for Monitoring Process Mean and Variability," *Quality Technology and Quantitative Management*, 1, 65-85.